

FermiCloud Testing and Development Platform Project Plan

Steven Timm

Version 3.2
Last modified 04-Sep-2009

Abstract:

This document describes the program of work for cloud computing testing and development at Fermilab.

Document Revision History:

Version	Date	Author	Comments
1.0	02-Jan-09	Steven Timm	First draft of project plan
2.0	21-Jul-09	Steven Timm	Second draft of project plan
3.0	17-Aug-09	Steven Timm	Version to be presented to CD Management
3.1	18-Aug-09	Steven Timm	After K. Chadwick revisions
3.2	04-Sep-09	Steven Timm	After G. Garzoglio and E. Berman revisions

Executive Summary:

This project will investigate, design, and deploy a flexible Infrastructure-as-a-service facility (“FermiCloud”) for use by the grid and storage developers, integrators, and testers of the Fermilab Computing Division. We are investigating currently-available virtualization and provisioning technologies, discussing requirements with stakeholders, and determining the budget in FTE’s and hardware costs. This project will deliver an infrastructure service which deploys just-in-time build and test images for all supported Fermilab operating systems. The virtual images will be operational only as long as they are needed. This will save on floor space and power, and also save significant administrative effort currently used to provision and support legacy hardware. The experience we gain for this project will also be used in the proposed General Physics Computing Facility currently under design for the Intensity Frontier experiments, as well as other stakeholders across the division. We are executing the project in three phases:

- Phase zero is the planning and requirements gathering phase of the project, detailed in this document.
- Phase one is the deployment of a small cluster on existing hardware for proof of principle.
- Phase two is the deployment of a cluster of two racks of new servers as requested in the FermiGrid Tactical Plan.

The scope of this project is to build a scientific testing and development private cloud at Fermilab. We will not deploy services on commercial clouds, although we will attempt to find a solution which is compatible with existing commercial clouds. Nor does the project include virtualization of the classic IT and business functions of the lab, which is under investigation by a different group and quadrant. This project does not include virtualizing the worker nodes of FermiGrid nor submission of virtual machines as jobs to FermiGrid.

Stakeholders:

FermiGrid Services Group (Timm, Chadwick)
Open Science Grid Group (Garzoglio et al)
Data Movement + Storage (Oleynik et al)
Fermilab Experimental Facilities dept. (J. Allen et al)
High Performance Parallel Comp. Fac. (Simone, Singh)
Comp. Enabling Technologies (Kowalkoski)
CMS Tier 1 (Holzman, Dumitrescu)
OSG + LHC Tier 3 Fraser (OSG), Snihur (CMS), Benjamin (ATLAS).
OSG Storage (Levshina)
REX Dept (D. Box)
SSE Group (Garren)

Statement of Need:

The initial need for infrastructure-as-a-service comes from the requirements of developers and integrators. The various development and integration groups within the SCF quadrant, including WAN,

Data Movement and Storage, Open Science Grid, and FermiGrid, rely on the FCC1 GridWorks cluster (a.k.a. FAPL cluster) for development and testing. Most of the machines in this cluster are 6 years old or older, and were not in good condition when they were made available to us after being decommissioned from the production farms. Over the past year almost 50% of these machines have broken and been unfixable, causing delays in critical development and testing work. In addition, there are a number of other development machines around FCC1 and FCC2 which are also aging and have no replacement plans at present. Even if working properly, the current machines have inadequate memory to test grid middleware. The developers and integrators need access to modern hardware, particularly for intense integration and testing cycles. This access is often needed on short notice. An on-demand provisioning system, coupled with the capacity to store the state of the system when it is not in use, is needed.

Since this system was proposed, discussions with major stakeholders and groups which manage large amounts of servers within the Division have identified virtual machine provisioning technology as a key technology in several new and ongoing activities. These include the following:

1. General Physics Computing Facility, the proposed home of interactive and small batch analysis for the Intensity Frontier experiments. Virtual machine technology will be used in deploying experiment-specific requirements. It will begin with static virtual machines such as are currently used by FEF or FermiGrid. The program of work of dynamic provisioning in this project will allow a transition to dynamic provisioning in the General Physics Computing Facility. This facility will also include virtual machines with identical configuration to FermiGrid worker nodes so that users can log in and debug their jobs.
2. Server OS Virtualization and provisioning technology. FEF and FGS rely on server virtualization to run large numbers of low-cpu-usage servers on a few machines. FEF has used Virtual Iron with success and is actively investigating its follow-on product, OracleVM. FGS has used open-source methods. Upcoming transitions in the commercial virtualization space and the open-source world make this a good time to evaluate potential new solutions.
3. Storage Evaluations--G. Garzoglio is currently leading an investigation of current storage use cases on Grid computing at Fermilab and has proposed a program of work which includes further testing of distributed file systems such as Hadoop and Lustre. Such file systems are often closely associated with cloud computing work and an infrastructure-as-a-service platform provides a good testbed to do this work with minimal impact on the sysadmins.
4. LQCD—They are considering distributing virtual machines to remote users for kerberos client and login, also considering the possibility of making 4-processor or 8-processor virtual machines available to users for compilation and testing of small jobs, with the idea that users would work some of the obvious problems out of their job before they get on the main cluster. They have used VirtualBox for their current work.
5. OSG and LHC—There is active effort in the LHC experiments to make a set of distributable virtual machines so that university sites can actively provision a Tier 3 with a minimum of effort. These typically can run under Xen or VMWare. The CERNVM project to create virtual CERN environments on the laptop and desktop is also a significant strategic direction for them.
6. External cloud investigations—Up until now these have included running production ([STAR@BNL](#)), investigation of using cloud for disaster failover (CMS). Several groups are interested in the external cloud as a development platform.
7. Virtual Services Group (LSCS quadrant) is virtualizing a number of servers using VMWare ESX.

Available Technology Investigation:

There are four basic areas of technology investigation identified thus far. (1) Available virtualization hypervisors and the hardware and OS on which they run. In common use at Fermilab thus far are Xen and VMware. RedHat's KVM is available in some test clusters. Microsoft's Hyper-V is also used in the Windows servers. Commercial as well as open source versions of Xen will be considered. (2) Available virtual machine provisioning and deployment mechanisms. These include open-source emulations of the Amazon Elastic Compute Cloud (EC2) such as Eucalyptus (UCSB) and Nimbus (U of Chi.) They also include Virtual Iron (soon to be incorporated into Oracle VM) and a number of other commercial offerings, including some that RedHat currently has in beta such as oVirt and OpenQRM. (3) Cloud-like object stores capable of storing and retrieving virtual machine input and other files. (4) Highly performant clustered file systems that work under Xen and the hardware to support them, needed for VM storage in (3) and also non-system-image data access. The project will attempt to use existing software systems to the greatest extent possible, but some minor modifications are expected to be necessary to comply with Fermilab's security environment. The initial goals of developer support can be met without a high-performance backend store if necessary.

Requirements:

The requirements list is not complete as yet, but a large number of requirements have been collected. We are aware of the following CD groups who have contemplated deploying virtualized systems in some form, or have actually done so. FermiGrid (Timm/Chadwick/Yocum), FEF (Allen), CMST1 (Holzman), CSI/WST (Fidler), LQCD+CET (Singh, Kowalkowski). We are aware of the following potential customers for a virtualized development cluster: OSG Group (Garzoglio), and Enstore/dCache developers (Moibenko/Perelmutov). These three groups are the major customers for the (barely) existing GridWorks cluster. Also having expressed interest are CDF, for various grid submit head nodes, D0 for various miscellaneous servers, ILC/Physics Tools (Garren), and CSI (soon-to-be-decommissioned build cluster), LQCD (virtual machines for build/test/debug plus small jobs). The General Physics Computing Facility envisioned for the Intensity Frontier experiments also plans to make use of these technologies.

A. Operating systems that need to be supported:

1. Client virtual machines: Linux (any versions currently within Fermilab baselines).
2. Server host machines: Linux (any versions currently within Fermilab baselines).
3. Desktop host machines: Linux, Windows, Mac.

B. Functionality:

1. (FGS) Ability to deploy worker node virtual machines to offer batch slots for opportunistic use on FermiGrid when developer demand is low.
2. (FGS) Emulation of the Amazon EC2 API as well as other evolving open-source API's

C. Estimated number of total virtual images needed by developers:

1. (FGS/OSG) Estimated demand within Grid department: OSG ReSS, 6 VM's, OSG Storage, 20-30 VM's, Grid Services related maintenance and testing 10 VM's (includes GlideinWMS), Fermilab Grid Storage Evaluation, 30-40 VM's, Gratia testing 12 VM's, SAMGrid maintenance and testing 4-6 VM's, OSG nightly install tests 10 VM's. Science dashboards 12-15 VM's. These are in addition to any virtual

- machines which are now available on the FermiGrid test nodes, which are virtualized but whose virtual machines are not dynamically provisioned, there are approximately 10 of these at the moment.
2. (DMS) 15-20 VM's for SRM and dCache deployment.
- D. Estimated number of images that need to be deployed simultaneously
1. About 50% of the above, 40-50 VM's.
- E. Typical memory and disk footprint of application
1. Varies significantly but average is between 1-2 GB of RAM.
- F. Potential disk and network performance issues due to virtualization.
1. (FGS) These are still under investigation but initial studies have shown that paravirtualized Xen, the most common solution, has performance very close to native disk and network access so we do not anticipate issues for the applications mentioned above.
- G. Configurability, maintainability, documentation
1. (FEF) Management GUI is essential, willing to buy commercial product to get it.
 2. (FEF) Amount of installation, configuration, customization on part of admins should be minimal.
 3. (FEF/FEF) Good monitoring functionality for machines that are already up.
- H. Network topology
1. (FGS) Need plan to serve either dynamic or fixed IP's. Latter are needed for grid services which expect a non-DHCP hostname.
 2. (FGS) Are private VLAN's needed on a customer-by-customer basis? Several other academic and commercial clouds have used this model successfully. Need to investigate.
 3. (FGS) We have requested and expect to get a dedicated subnet for these investigations.
- I. Security and patching
1. (FGS) Need a patching mechanism to fit the OSE baseline, and not let the virtual machine onto the public net until it is patched.
- J. VM and OS provisioning requirements
1. (FGS) Virtual machine provisioning and OS provisioning should be completely automated and transferable to groups which manage large quantities of worker nodes.
 2. (FGS/OSG) Virtual machines should be able to be paused and saved and restored later on other hardware.
 3. (FGS) Virtual machines should be able to be turned on and off at a fixed time, and also be able to be instantiated to go for a certain length of time.
 4. (FGS/FEF) Important to have live migration and failover of virtual machines in case of hardware failure or scheduled maintenance

Design and Evaluation:

Based on the requirements collected above, we will conduct the investigations and come up with a design for a full system that meets them. Our current estimates for the FTE effort and equipment costs are described in the FermiGrid Tactical Plan and the Grid Services tactical plan. We have already deployed eight worker nodes using static virtualization techniques in GCC Room B. These are currently serving as the OSG Persistent Integration Test Bed. By use of virtualization we

have been able to deploy a Condor, PBS and SGE cluster for Integration use in the same space as was currently used only for a Condor cluster before. Within the next month we will bring 16 more Dell Poweredge 1950 nodes into this test cluster and start the dynamic provisioning techniques. These nodes were all initially bought for testing purposes in FY2007. Once a suitable set of provisioning techniques is identified, we will present them and proceed with acquisition of the two racks of nodes for Phase Two, as specified in the FermiGrid Tactical Plan.

WAN Group Network Testing:

Due to the high bandwidth requirements of the WAN group testing and the experimental kernel work that they do, some of the newer GridWorks cluster nodes on FCC1 and the existing switch gear will need to remain available for their testing so they can run on bare metal. Their testing is not a good candidate for virtualization and is not intended to be integrated into FermiCloud.

Deliverables:

1. Test Cloud Cluster up for investigations: Late September 2009
2. Design plan for production FermiCloud cluster, including finalized hardware and software technical requirements and Statement of Work: Late October 2009
3. Available Technology Investigation report complete: January 2010
4. Requirements collection from potential VM users complete: February 2010
5. Production Cloud VM provisioning, OS Provisioning, and security patching solution in production: Late March 2010.
6. Generalized VM library for users of pre-built virtual machines, late March 2010.
7. Transfer of all OSG Group development and integration work and DMS storage development work to FermiCloud—April 2010 or sooner if existing FAPL cluster is forced to move off FCC1 before that time.
8. Integration of FermiCloud technology into other facilities such as General Physics Computing Facility—June-July of 2010.